


Report from  
HEPiX/HEPNT Fall '03

Stephan Wiesand  
DESY - DV -  
Nov. 24<sup>th</sup>, 2003



## This is not a LISA '03 report


LISA '03

OCTOBER 26-31, SAN DIEGO, CA

BY & FOR SYSADMINS!

THE 17TH LARGE INSTALLATION  
SYSTEMS ADMINISTRATION CONFERENCE

But there is some  
additional input on  
HEPiX topics from  
LISA



DESY participants: W. Friebel, P.v.d.Reest, S. Wiesand

# HEPiX



- politically correct name: **HEPiX/HEPNT**
  - meeting held twice a year, in spring and fall
  - on either side of atlantic ocean every other time
- **Fall '03** held at **TRIUMF** in **Vancouver** (organized very well)
  - DESY participants (73 total)
    - **R. Baltrusch, P. v. d. Reest**
    - **W. Friebel, H. Schwendicke, Stephan Wiesand**
  - **<http://www.triumf.ca/hepix2003/>**
    - all **presentations** (**powerpoint** or **staroffice** or **PDF** format)
    - **audio/video capture** of almost all sessions (actually very usable)
    - **good summary** (23 pages of text in PDF format)

# HEPiX Fall '03: Format



- 3 days of **site reports** and **general talks** (many very good)
- 1 day dedicated to presentations on **security** (ditto)
- $\frac{1}{2}$  day of parallel sessions
  - **security round table**
  - **windows round table**
    - about security, first part joint security session
    - **mass storage forum** (not covered in this talk, P.v.d.R. only)
- first HEPiX with commercial **vendor demos**
  - demos and talks by 2 vendors of **advanced, global file systems**
- **invited talks** by Red Hat & Microsoft

# Format of this Presentation



- good summary exists
  - => no point in reproducing each talk in three sentences
- instead, focus on
  - selected topics from site reports and general talks
    - by topic
    - Windows input by R. Baltrusch, H. Schwendicke
    - mail/spam input by W. Friebel
  - security
  - linux distribution discussion
  - providing additional information not available on the web



# Omissions

- many interesting topics are not covered here:
  - **GRID** (rollout, security)
  - **Fabric Management Tools**
    - RPM updates, Host Databases, Windows Remote Installation, ...
  - **Storage: Castor, DCache, Backup, ...**
  - **PKI** (good tutorial)
  - **Windows NT -> AD migration**
  - **AFS cross cell authentication**
  - **Unix -> Windows migration (MS SFU, ...)**
  - **Services (CVS, Solaris, General Email, Printing, ...)**
  - ....

# Topics from Site Reports: OSs



- Operating Systems
  - Linux, Windows, Solaris/SPARC everywhere
    - some HP-UX, AIX, IRIX left, typically being phased out
    - little MacOS (X) support, typically not on agenda
    - Windows rules the desktop domain
    - Linux rules the compute server domain
    - Linux is conquering the "real services" domain at many sites
      - AFS, NFS, Oracle, TSM, ...
      - mail, DHCP, Web, DNS, ...
  - all sites **concerned** about **linux distributions**
    - some expressed interest in Solaris/x86
      - SUN was marketing it very actively at LISA

# Site Report Topics: Hardware



- **Complaints** about P4/**Xeon**
  - **performance/GHz** much worse than PIII
    - HyperThreading helps, but issues with linux scheduler, and CPU accounting / job scheduling complicated
  - **power consumption**
    - "Westgrid" at UBC (1008 dual Xeon 3GHz) can not run all blades (IBM bladecenters) in a crate until power supplies replaced
- positive reports about AMD **Opteron** performance
  - being **considered** for most farm purchases next year
- one site reports **SCSI-attached IDE-RAID** was a **desaster**



# Hardware continued

- CERN seems last site settling for "white boxes" for farms
  - and pursuing custom low cost solution for serial console access
- most others are back to "professional" systems
  - rack mounted
  - sometimes with KVM switches
  - some sites require remote power control
  - some sites require management/monitoring software
- one site reports power supply quality is an issue
  - bad ones cause severe distortion of grid voltage
    - -> high peak currents

# Site Report Topics: Windows



- all sites have or are deploying **AD domains**
  - 2000, 2003, XP
  - NT/9x still exist at some sites
- most sites have deployed or are evaluating at least one of
  - **MS SMS**
    - **systems management server**
  - **MS SUS**
    - **software update service**
  - necessity for efficient **patch deployment**
  - **typically, only for new domains**
    - NT/9x often managed manually only

# Site Report Topics continued



- Windows Terminal Services
  - either already deployed
    - sites report use increasing
    - often citrix
  - or being evaluated (most other sites)
    - typically RDP
- SLAC project on AD/Heimdal password synchronization
  - working with MS on tools to allow this smoothly
  - interest expressed by DESY Windows group
- Kerberos 5 is present or most likely future at all (?) sites
  - desire for single sign on expressed by some



# Selected Topics: SPAM

- all labs have to throw significant resources on spam fighting
  - first thoughts about a **whitelisting** approach (CERN)
    - mails from whitelisted sources pass
    - automatic response to non whitelisted sources
    - response to this mail passes
      - likely to happen for real mail only
      - most spam has no valid from: or reply-to: address
    - user can then whitelist the source if desired
    - **problems:**
      - **hassle** for senders & users
      - **what if both sites do this ?**
  - most postmasters seem not to like the idea

# Selected Topics: WebDAV



- extension to http protocol
  - potentially: https (but that's the future)
  - available with IIS, Apache, ...
- native filesystem drivers exist for Windows, OsX, Linux
- could be a solution for file sharing between these OSs
  - securely, over wide area networks
- CERN running a pilot gateway to their (MS) DFS
- looks promising, [linux davfs seems dead to me though]
- many limitations in practice today
  - no SSL, port 80 only, authentication methods, ...



# Security

- most major labs had a high ranking security officer present
- security officers at all sites had an "interesting" year
  - Windows worms & viruses
    - Slammer, Sobig, Lovsan, Welch,...
    - temporarily caused up to 30% packet loss on internet
    - effectively shut down some labs (and enterprises)
    - infected systems within minutes
      - during (re-)installation
      - before systems could be patched when turned on
    - CERN hit by virus before antivirus signature available
      - exploits IE weakness, installs spam relay on random high port
      - lab faced threat of being brought to court due to nature of spam



# Security continued

- Linux **ptrace** vulnerability
  - trivially exploited from cracked user accounts
    - success rate almost 100%, exploit widely available
- frightening **root kits**, like SuckIt
  - very good at concealing itself, very hard to detect
  - installs backdoor defeating all firewalling
    - listens on ALL ports for backdoor trigger packets
    - then initiates TCP connection from infected host
- **users** running
  - P2P filesharing software
  - IRC (and being caught by bots)
  - vulnerable sshd or httpd or... (on high ports)

# Security: Common Problems



- common agreement today these are the worst problems:
  - systems **not** properly (professionally) **managed**
    - each of these measures alone almost eliminates attack potential:
      - applying **patches** timely
      - running **antivirus** software with daily updated signatures
      - running a **personal firewall** at least buys time
    - how could so many systems be compromised this year ?
      - **fix for many attacks available weeks / months / years before !**
  - **firewall penetration**
    - **notebooks, VPN, dialup (home systems)**
    - **unauthorized, vulnerable services / applications**
  - users downloading **malware**, opening unknown attachments, ...
  - notebooks that can only be updated inside their home network
    - one week can be too long these days



# Security: Common Measures

- most sites now apply these or are planning to do so:
  - all **devices** attached to network **must be registered**
  - and **responsible has to agree (in writing) to rules**, like
    - system must be configured securely
    - patches must be applied timely, system rebooted if necessary
    - system must be running update antivirus and firewall
    - system must not be running unauthorized services
  - **users** of centrally managed systems **must agree to rules**, like
    - no P2P software or other unauthorized services / applications
  - **VPN/dialup users must agree to rules**, like
    - no additional software, no usage by the kids, ...

# Security Measures: Exceptions



- exceptions from rules generally granted if necessary
  - if work cannot be done without violating the rules
- most sites require a written statement
  - why there's a need for it
  - what technical measures prevent security breaches
    - "how will you prevent unauthorized file access through your P2P filesharing application?"
  - signed by user and responsible
- sites report almost all requests are withdrawn after pointing out this requirement



# Security Measures: Scans

- major sites run scans of their network
  - detect vulnerable systems, unauthorized services
  - detect compromised systems (backdoors, ...)
  - full scans regularly
    - typically take  $O(1 \text{ month})$  to complete
  - individual scans immediately when new devices attached
  - problems:
    - scan may disrupt operation of some devices (DAQ equipment...)
      - -> first detect OS, then apply specific scan
    - feasible to quarantine new systems until scanned?
  - vulnerable/compromised systems disabled on network level

# Security Round Table Results



- HEPiX labs will agree on **common set of minimal rules** for systems to be attached to their networks
  - systems carried by guests from other HEP labs are expected to comply with these (this is **YOUR notebook**)
- incidents and attacks should be communicated to the (closed) security mailing list
- a **new security discussion list** for HEPiX was created
  - not public, but **open** to anyone from any HEP lab
    - subscription must be approved by list owners (hosted at fermilab)
    - **new members are expected to introduce** themselves
      - or may be removed from list



# Security: Summary

- today's threats are serious
  - no major damage yet, but only matter of time
- "patch early, patch often!"
  - any system, centrally managed or not
  - including network gear, farms, desktops, notebooks, ...
  - this is a significant departure from
    - "choose patch time wisely for optimal availability"
    - "it's ok to patch servers only"
    - "locally only exploitable bugs aren't worth patching"
- firewalls can help, but are **not** a sufficient solution
  - limit exceptions as much as possible

# NB: related talks at LISA



- San Diego Supercomputing Centre:
  - they have no firewall at all
    - "we have to secure all nodes anyway"
    - "=> resources better spent on nodes than on a firewall"
  - lively discussion, no real objections
- Argonne National Lab:
  - how they made their lab secure
    - without spoiling the scientist's work
  - significant effort
  - no chance of success if management doesn't play along
    - incentive: DOE labs have to pass security audits, or budget suffers...

# The Linux Discussion: Background



- almost all HEP sites run some vanilla **Red Hat Linux**
  - many also already run a few Red Hat Enterprise Servers
    - typically for Oracle
    - significant cost per server and year
- some (DESY, GSI) run SuSE and/or debian
  - few SuSE/debian hosts at few other sites
- Red Hat early this year shortened **distribution life times**
  - to 12 months
- later this year they **discontinued** their vanilla distribution
  - superseded by **Fedora**, life time 6-9 months

# Linux Discussion: Background



- distribution **end of life**:
  - RedHat 7.x      **12/03**
  - RedHat 8.0      **12/03**
  - RedHat 9        **04/04**
  - Fedora Core 1   **07/04**      (at best, and limited)
  - SuSE 8.2        **04/05**
  - SuSE 9.0        **10/05**
  - debian woody    **12/04 + ?**      (12 months after undefined date)
- SuSE/Red Hat **Enterprise distributions live 5 years**
  - = unlimited in practice

# HEPiX Linux Discussion



- most labs now have to find a new workhorse distro soon
  - CERN & probably others will support 7.3 until 12/04
  - but need several months for certification of new OS
- most labs have contacted distributors about volume licensing
  - we talked to SuSE and RedHat, all others to Red Hat only
  - all got similar offers around XXX \$/year/node
  - no lab could negotiate acceptable conditions so far
- => try common HEP effort
  - Red Hat invited to HEPiX
  - session on this topic (w/o RedHat presence, w/o recording)

# Red Hat at HEPiX



- Red Hat sent Don Langley
  - sales manager for california
    - including SLAC
- held a **plain marketing** talk for Red Hat Enterprise Linux 3
  - session not recorded
  - pdf on the web
  - no additional information
- **refused to discuss** HEP volume licensing
  - just stated they're "**interested in creating a win-win situation**"

# Summary of Discussion Session



- most sites really want to use Red Hat Enterprise Linux
  - debian/SuSE/others not considered seriously
- but **not** with their default support model
  - HEP sites most of all want the patches
    - not per incident remedial services
    - after inserting an own kernel module, these are void anyway
    - on LISA, heard complaints about service from people having it
  - some sites interested in RHN satellites (->delegation)
- HEPiX believes Red Hat have not yet made up their mind
  - give them more time (how much ?)
- try negotiating on higher level

# Other Linux Options discussed



- some consider rebuilding a RHEL from sanitized source
  - after all, it's GPL
  - probably legal if all trademarks and files with other licenses (artwork) are removed, and the name is changed
    - situation is not really understood by anyone
    - CERN would require written permission before redistribution
- some consider using Fedora
  - and hoping for Fedora Legacy to work
    - volunteer project hosted by Red Hat to provide patches for old fedora
- hardware vendors may offer reasonable RH WS licenses
  - but can this be extended to existing hardware?



# Linux in HEP: Next Steps

- CERN, SLAC, Fermilab will try to negotiate with Red Hat
  - objective: acceptable conditions for using RHEL
    - in all HEP (LCG?) labs, and collaborating institutes
  - no deadline set
- U.S. department of energy is negotiating for all their labs
  - what if they succeed, and HEP doesn't ?
- DESY will watch from the side line
  - we're about to roll out DL5 based on SuSE 8.2
    - buys us a year, no immediate pressure
  - but we expressed interest in buying into a reasonable solution



# HEPiX Aftermath

- negotiations are going on
  - state of discussion is not being disclosed
    - except for a claim (Nov. 21<sup>st</sup>) that there is hope for
      - a satisfactory solution for all of HEP
      - "within a few weeks"
- White Box Linux became available (rebuilt RHEL source)
  - some problems
    - getting rid of all artwork is a significant effort
    - some source RPMs as downloaded do not yield working packages
  - at least one HEP lab is doing the same now
  - Red Hat have no reason to make this easy...



# Conclusion

- **HEPiX/HEPNT** is very **relevant** to DESY computing
  - many good talks could not be covered here
    - have a look at them on the web
- **DESY** staff **should attend** regularly
  - traditionally, intentionally **inexpensive**
- next meetings:
  - May 2004 in Edinburgh
  - Fall 2004 at Brookhaven